

The Plateau Problem in the Heteroskedastic Probit Model*

Eric Freeman[†] Luke Keele[‡] David Park[§] Julia Salzman[¶]
 Brendan Weickert^{||}

August 14, 2015

Abstract

In parameter determination for the heteroskedastic probit model, both in simulated data and in actual data collected by Alvarez and Brehm (1995), we observe a failure of traditional local search methods to converge consistently to a single parameter vector, in contrast to the typical situation for the regular probit model. We identify features of the heteroskedastic probit log likelihood function that we argue tend to lead to this failure, and suggest ways to amend the local search methods to remedy the problem.

*We thank Michael Alvarez and John Brehm for kindly providing us with their data set. For comments and suggestions, we thank Chris Achen, Neal Beck, John Brehm, Andrew Gelman, Gary King, Andrew Martin, Walter Mebane, Kevin Quinn, Jasjeet Sekhon, Jim Stimson, Marco Steenbergen, Greg Wawro, Richard Williams, Fred Boehmke, and the Statistics Working Group at Columbia University. We thank Andrew Gelman for making this collaboration possible by introducing three of the authors.

[†]Institute for Research on Labor and Employment, University of California at Berkeley

[‡]Department of Political Science, Penn State University, Corresponding Author

[§]Applied Statistics Center, Columbia University

[¶]Department of Biochemistry and Department of Statistics, Stanford University

^{||}Palantir Technologies

A common difficulty encountered in the estimation of statistical models is that of unequal variances or heteroskedasticity. In the context of ordinary least squares, heteroskedasticity does not bias parameter estimates; rather, it either inflates or underestimates the standard errors. Heteroskedasticity, however, is more problematic in discrete choice models such as logit or probit and their ordered and multinomial variants. If we have nonconstant variances in the error term of a discrete choice model, not only are the standard errors incorrect, but the parameters are also biased and inconsistent (Yatchew and Griliches 1985).

Alvarez and Brehm (1995) generalize a basic model for heteroskedasticity developed by Harvey (1976). They call these heterogenous choice models, which include heteroskedastic probit and heteroskedastic ordered probit models to correct for unequal variances with discrete outcomes. They also used heteroskedasticity as a means of exploring heterogeneity in choice situations. These heteroskedastic probit models have been frequently used to explore heterogenous choices and behaviors (Alvarez and Brehm 1997, 1998, 2002; Busch and Reinhardt 1999; Gabel 1998; Lee 2002; Krutz 2005). Routines for these models have become standard in statistical software such as **Stata**, **Limdep**, **SAS**, **Eviews** and **Shazaam**.

We use simulations to explore the optimization difficulties that can arise during the estimation of these models. In our simulations, the models are correct. Thus, the anomalies that we find cannot be explained by specification error. We argue that estimation difficulties are due to the functional form of these models. In this paper, we observe the following dichotomy: the usual algorithms for optimization of regular probit log likelihood often fail with the heteroskedastic probit model. It is likely that the problems with these model are under-reported in the literature, since the errors that are encountered may give rise to the “file-drawer” problem where difficulties with these models go unreported as these papers are sitting in investigators’ file drawers (Iyengar and Greenhouse 1988). The point of this paper is, first, to give evidence of that dichotomy, second, to attempt to identify those features of the heteroskedastic likelihood function that lead to the failure, and third to suggest ways a local search might be adapted to those features.

To this end, in section 1 of this paper, we give a brief exposition of the regular and het-

eroskedastic probit models. In section 2, we investigate the performance of “out-of-the-box” optimization techniques with the heteroskedastic probit model, first through simulated data, and second through a data set collected by Alvarez and Brehm (1995). In section 3, we delve more deeply into our simulated data to attempt to gain some intuition about the graph of the likelihood function, and thus about the failure of local search methods, in the heteroskedastic case in general. And in section 4, we suggest a modification to tailor optimization methods to the heteroskedastic probit model, and offer additional recommendations.

1 A Probit Model with Heteroskedasticity

We first define the model, beginning with the standard probit model. Fix a constant k_1 –dimensional vector β , and assume that, given an observed value \mathbf{x}_i of a k_1 –dimensional vector of random variables and an observed value ϵ_i of a random variable distributed normally with mean 0 and variance σ^2 , a latent random variable y^* takes the value¹

$$y_i^* = \mathbf{x}_i' \beta + \epsilon_i.$$

This value is not seen by the researcher. A researcher observes only the value

$$y_i = \begin{cases} 1 & \text{if } y_i^* \geq 0 \\ 0 & \text{if } y_i^* < 0. \end{cases}$$

One then readily calculates that

$$\Pr(y_i = 1 \mid \mathbf{x}_i) = \Phi \left(\frac{\mathbf{x}_i' \beta}{\sigma} \right),$$

¹To be completely precise, we note that we think of vectors as column vectors.

where Φ is the cumulative distribution function for the standard normal distribution. Only the parameter $\frac{1}{\sigma}\beta$ is identified, so typically one assumes

$$\sigma = 1$$

for purposes of identification.

In the heteroskedastic probit model, one allows σ to take values other than 1, by estimating a model for its value. This may arise in many situations, for example, in a model of personal choice where levels of information vary across individuals. A convenient form for modeling σ is

$$\sigma = \exp(\mathbf{z}'\boldsymbol{\gamma}),$$

where \mathbf{z} is a k_2 -dimensional random vector and $\boldsymbol{\gamma}$ is a k_2 -dimensional parameter vector. Thus the model can be written as

$$\Pr(y_i = 1) = \Phi\left(\frac{\mathbf{x}_i'\boldsymbol{\beta}}{\exp(\mathbf{z}_i'\boldsymbol{\gamma})}\right).$$

A natural way to think of this model, as above, is as consisting of a latent variable given by

$$y^* = \mathbf{x}'\boldsymbol{\beta} + \epsilon,$$

where

$$\epsilon \mid \mathbf{z} \sim \mathcal{N}(0, \exp(2\mathbf{z}'\boldsymbol{\gamma})). \quad (1)$$

But one can also think of it as a model with a latent variable with mean $\frac{\mathbf{x}'\boldsymbol{\beta}}{\exp(\mathbf{z}'\boldsymbol{\gamma})}$ and a disturbance term distributed as a standard normal variable. We also observe that there can be conditions in which the model is not identified, for example if \mathbf{z} consists only of a constant.

One can estimate this model using maximum likelihood, much as one does for the probit model. Assume that there are n observations of y_i , \mathbf{x}_i and \mathbf{z}_i . For convenience, let \mathbf{y} be the

$n \times 1$ vector of all observations of y_i , let X be the $n \times k_1$ matrix whose i^{th} row is \mathbf{x}'_i and let Z be the $n \times k_2$ matrix with i^{th} row given by \mathbf{z}'_i . Then the log likelihood function is given by:

$$\ell(\boldsymbol{\beta}, \boldsymbol{\gamma} \mid \mathbf{y}, X, Z) = \sum_{i=1}^n \left(y_i \ln \Phi \left(\frac{\mathbf{x}'_i \boldsymbol{\beta}}{\exp(\mathbf{z}'_i \boldsymbol{\gamma})} \right) + (1 - y_i) \ln \left(1 - \Phi \left(\frac{\mathbf{x}'_i \boldsymbol{\beta}}{\exp(\mathbf{z}'_i \boldsymbol{\gamma})} \right) \right) \right). \quad (2)$$

Given a set of n observed values of the random vectors y, \mathbf{x} and \mathbf{z} , in order to obtain the maximum likelihood estimate, one maximizes this function over the space of possible choices of $(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathbb{R}^{k_1+k_2}$.

2 Search Algorithm Performance

We now examine how standard optimization algorithms perform when maximizing the heteroskedastic probit log likelihood function.

2.1 Simulated Data: Heteroskedastic Probit Model

We first consider the behavior of search techniques when employed on simulated data sets. We note that this and other data sets used in the paper are not designed to be pathological. In this section, we describe the simulated data generating process and then use four different search methods on this one data set.

The basic data generating process is as follows. First, we generate 1000 observations of \mathbf{x}_i and \mathbf{z}_i , with components equal to 0 or 1. Then we generate a particular choice for both $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ that we call the *model parameters*, which we denote by $\boldsymbol{\beta}_0$ and $\boldsymbol{\gamma}_0$. Given these choices, for each observation i , we generate a disturbance term ϵ_i sampled from the normal distribution with mean 0 and variance $\exp(\mathbf{z}'_i \boldsymbol{\gamma}_0)$. Then we consider each sum

$$\mathbf{x}'_i \boldsymbol{\beta}_0 + \epsilon_i. \quad (3)$$

If this sum is greater than or equal to 0, we define y_i to be 1, and we define y_i to be 0 if the

sum is less than 0. In this fashion, we generate the parameters β_0 and γ_0 , and the data set $\{\mathbf{x}_i, \mathbf{z}_i, y_i ; 1 \leq i \leq 1000\}$.

Now consider one potential complication. Suppose that for all i , the term $\mathbf{z}_i' \gamma_0$ is large and negative, so that the variance $\exp(\mathbf{z}_i' \gamma_0)$ is quite small. In this case, it could be that the term $\mathbf{x}_i' \beta_0$ in the sum (3) above is much larger in magnitude than that of the term ϵ_i for all observations, and thus may completely determine the value of y_i . We refer to this as a situation in which there is no *crossover*. Now define $\mathbb{1}(v) = 1$ for $v \geq 0$ and 0 for $v < 0$. Then, more formally, we say that there is crossover for observation i if

$$y_i \neq \mathbb{1}(\mathbf{x}_i' \beta_0),$$

that is, if the disturbance term leads to y_i having a different value from that which the linear term $\mathbf{x}_i' \beta_0$ alone would predict. If all of the observations exhibit no crossover, then it seems intuitively likely that there will be no optimal choice for γ ; the basic idea is the following.² Suppose, to the contrary, that there is an optimizing argument and suppose for example that it occurs when $\beta = \beta_0$. For those observations with either $\mathbf{x}_i' \beta_0 > 0$ or $\mathbf{x}_i' \beta_0 < 0$, because the components of \mathbf{z} are nonnegative, one can observe from Equation (2) that the summands of the log likelihood function will generally increase as the components of γ grow large and negative.³ For this reason, and more importantly to focus on the class of models for which the variance portion of the model plays an important role, we ensure that in the data generating process there is a significant percentage of crossover.

We now turn to the details of simulating the data. We take the dimension of \mathbf{x} to be three and of \mathbf{z} to be two. The first component of \mathbf{x} is a constant equal to one, while the second and third components, as well as the two components of \mathbf{z} , are realizations of Bernoulli variables with probability 1/2, i.e. discrete uniform $\{0, 1\}$ random variables. We

²This is of course exactly similar to the complete separation problem for the probit model. For one discussion of the complete separation problem, see Albert and Anderson (1984). Separately, we note that our purpose here is not to provide a complete proof, but only to explain the reasoning behind the choices in our data generating process.

³We assume here for simplicity that many of the components of the \mathbf{z}_i terms are strictly positive.

generate 1000 such observations \mathbf{x}_i and \mathbf{z}_i . As mentioned above, we refer to the parameters used in the data generating process as the *model parameters* and denote them by β_0 and γ_0 . The components of these parameters are themselves sampled from a (continuous) uniform distribution on $[-5, 5]$, although there are important provisos to this statement. For one, we select the second and third component of β_0 from $[-5, 5]$, but after doing so we in fact choose the first component so that the mean of $\mathbf{x}'_i\beta_0$ over all observations is 0. The purpose of this is to avoid the situation in which $\mathbf{x}'_i\beta_0$ is, say, positive for all or most of the data set; unless the variance is fairly large, we might expect to see very little crossover. The second caveat to the statement that the parameters are sampled from a uniform distribution is that we discard any choices (β_0, γ_0) of the parameters for which the crossover is less than 20% or greater than 30%; i.e., we simply repeat the random sampling until we find such a choice of parameters.⁴ Having selected β_0 and γ_0 , then for every observation indexed by i , we sample a disturbance term ϵ_i from a (pseudo-)normal distribution with mean 0 and variance $\exp(\mathbf{z}'_i\gamma_0)$. Then we define

$$y_i = \mathbb{1}(\mathbf{x}'_i\beta_0 + \epsilon_i).$$

This gives the complete set of simulated data $\{\mathbf{x}_i, \mathbf{z}_i, y_i ; 1 \leq i \leq 1000\}$. Now we fix one data set simulated in this fashion, and consider the performance of different algorithms when maximizing the log likelihood on this data set.⁵

We first present results for the algorithm commonly referred to as the BFGS algorithm (Broyden 1970; Fletcher 1970; Fletcher and Reeves 1964; Shanno 1970). In each of 1000

⁴The percentages chosen here are somewhat arbitrary, but reflect a desire for the average summand in the log likelihood function to be somewhat better (less negative) but not too much better than what turns out to be an essentially minimal value of $\log(1/2)$. If there is too much crossover, this average summand (which we later refer to as the normalized value of the log likelihood function) will generally be very close to $\log(1/2)$, whereas if there is very little crossover, it will generally be much better. In the former case, the value of the log likelihood at the model parameters is not much different from its value at what we call below the plateau solution, so it is hard to demonstrate the difference between the two. In the latter case, the model approaches a condition much like the case of complete separation in the probit model, so we want to avoid this situation. Recall as well that our purpose in this paper is merely to demonstrate that under reasonable conditions, there are concerns with optimization when using the heteroskedastic probit model, as opposed to showing that there are concerns under any conditions. Also, we note that we discuss why the value $\log(1/2)$ arises more below.

⁵The model parameters for this data set, to two decimal places, are given by $\beta_0 = (.56, -.31, -.84)$ and $\gamma_0 = (3.43, -4.07)$.

runs, we randomly select a starting value for the parameter (β_0, γ_0) to be estimated. Each choice of initial values has components in the interval $[-5, 5]$, so that each choice of initial values is in the box $[-5, 5]^5$. Denote the estimate resulting from the i^{th} run by

$$\left(\hat{\beta}_{BFGS,i}, \hat{\gamma}_{BFGS,i}\right),$$

and the set of all such estimates by

$$E_{BFGS} = \left\{ \left(\hat{\beta}_{BFGS,i}, \hat{\gamma}_{BFGS,i}\right), \quad 1 \leq i \leq 1000 \right\}.$$

We are interested in how close the estimates are to the model parameters. We employ two measures to measure the proximity of the estimates in E_{BFGS} to the model parameter values. The first measure is the simple Euclidean distance between each estimated parameter value and the model parameter value. The second measure starts by looking at the difference between the value of the log likelihood function F at the model parameter vector and the value of the log likelihood function F at each estimated parameter vector. We divide this difference by the number of observations in the data set. The idea here is to normalize the difference in order to better compare results using this data set with later results using a real data set, which has a different sample size. Here the number of observations is 1000, which we denote by n_{SIM} for expositional purposes.

We present histograms of these two measures over all 1000 runs. We look at two histograms, for frequencies observed in each of the following two sets, which correspond to the two measures described above:

$$D_{BFGS} = \left\{ \left\| (\hat{\beta}_{BFGS,i}, \hat{\gamma}_{BFGS,i}) - (\beta_0, \gamma_0) \right\|_2, \quad 1 \leq i \leq 1000 \right\},$$

where $\| \cdot \|_2$ denotes Euclidean distance, and

$$V_{BFGS} = \left\{ \frac{1}{n_{SIM}} \left[F((\beta_0, \gamma_0)) - F((\hat{\beta}_{BFGS,i}, \hat{\gamma}_{BFGS,i})) \right], \quad 1 \leq i \leq 1000 \right\}.$$

We note that $\frac{1}{n_{SIM}}F((\beta_0, \gamma_0)) = -0.45946$. There is in fact an added complication for the histogram for V_{BFGS} , which we discuss momentarily.

The two histograms are displayed in Figure 1. The left histogram plots frequencies for the set D_{BFGS} , and the right histogram plots frequencies for the set V_{BFGS} . The left plot uses a log scale on the horizontal axis, to more clearly present key phenomena. The right plot also uses a log scale on the horizontal axis, but only in the right section of the graph. There are actually almost 300 choices of starting values for which the log likelihood values at the estimates found using the BFGS algorithm are actually greater (i.e., less negative) than the value at the model parameters. This is not that surprising— we are using a finite sample, and obviously not a population. For the log likelihood function, which is the empirical analogue of the expectation of a generic summand over the population, we do not have any reason to believe that the *log likelihood* should be maximized at the model parameters; on the other hand, under suitable conditions we might expect that the *expectation* would be maximized at the model parameters.⁶

Naturally, these negative differences would not show up on the graph if we were to use a log scale, given that the log scale would only allow representation of those observations for which the log likelihood has a value at the estimated parameters that is in fact strictly less than the value of the log likelihood at the model parameters. We represent the observations for which the value of the log likelihood at the estimated parameters is actually greater than the value of the log likelihood at the model parameters in the graph in the following way. We plot a rectangle on the left side of the plot with height equal to the number of these cases and labeled as “Values Better than at Model Parameters.”⁷ As a final note concerning the two plots in Figure 1, observe that the scales for the horizontal and vertical axes differ by plot.

When starting with random initial values, the parameter estimates do not generally appear to be close to the model parameters, for about 70% of the 1000 runs. In fact, the

⁶See, for example, Lemma 14.1 of Ruud (2000), for a statement of a result of this type concerning the expectation.

⁷Its location along the horizontal axis is of course not meaningful.




Figure 1: Performance of the BFGS Search Algorithm on a Simulated Data Set for the Heteroskedastic Probit Model, for 1000 Random Choices of Initial Values

Euclidean distance between the parameter estimates and the model parameters is at least $3.16 (= 10^{-5})$ for all but 15 of the 1000 runs and at least 10 for 790 of the runs. Moreover, for more than 600 runs, the value of the log likelihood function at the estimated parameter values is at least 223.87 less than the value at the model parameters, which is -459.46 .⁸ Thus by two measures the BFGS algorithm appears to only perform well for some fraction

⁸The value .22387 equals $10^{-0.65}$, which is the left endpoint of the rightmost histogram cell in the plot on the right in Figure 1; recall that this value was obtained by normalizing, i.e. by dividing by 1000.

of starting values chosen randomly from a box, at least for this data set.

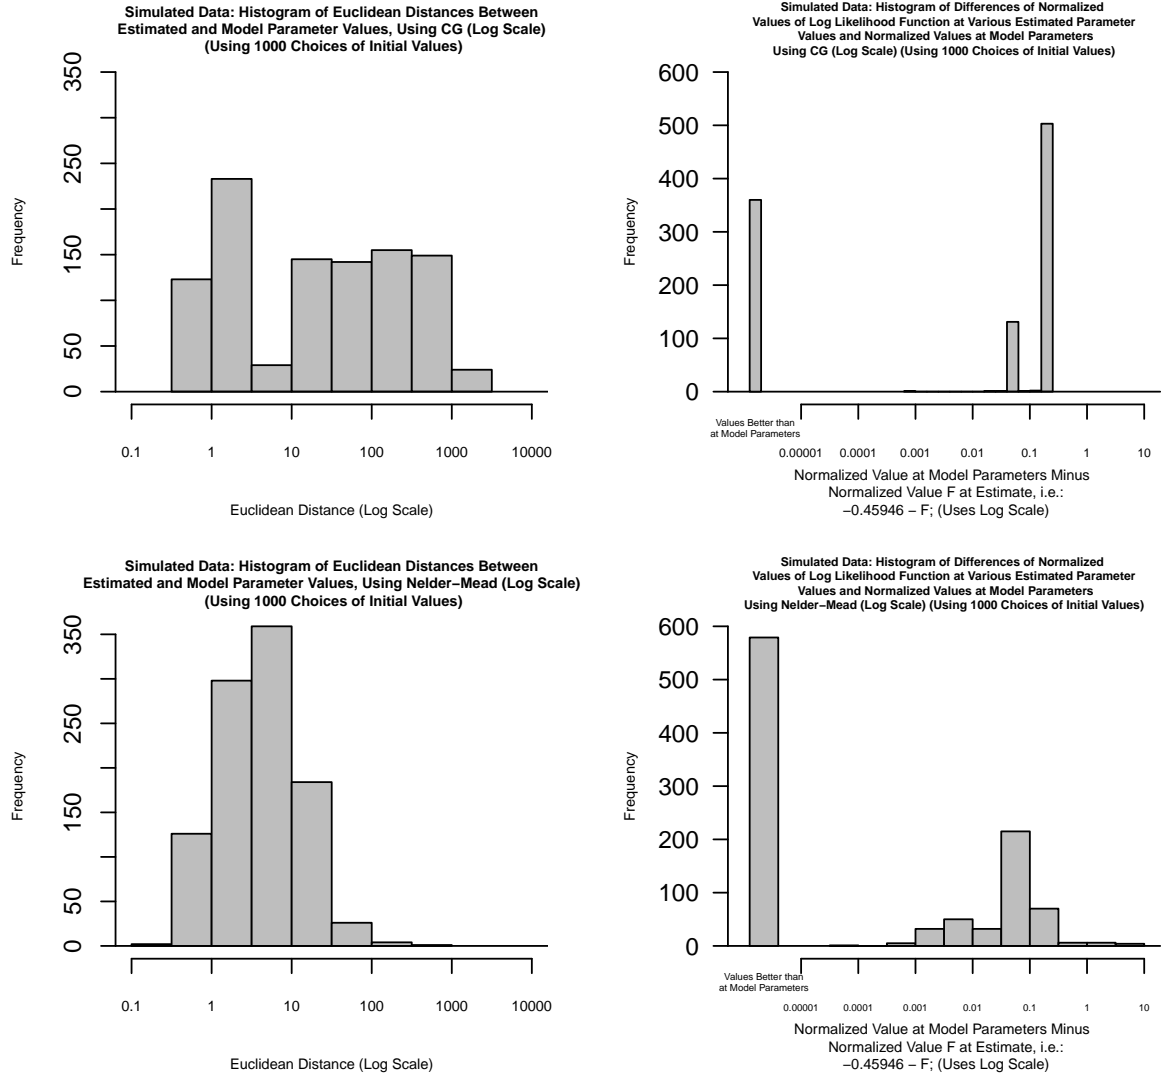


Figure 2: Performance of the CG and Nelder-Mead Search Algorithms on a Simulated Data Set for the Heteroskedastic Probit Model, for 1000 Random Choices of Initial Values

Figure 2 presents similar plots for the same simulated data set, but with two other optimization techniques. The simulations were run so that the same 1000 choices of initial parameter vector were made for each search technique. (To be clear, these are also the same initial choices used for the BFGS search technique.) We combine results for two separate algorithms in one figure. The top two plots present measures of the quality of estimates

found from implementing a conjugate gradients method (referred to as CG) developed by Fletcher and Reeves (1964), while the bottom two plots display results found using a method developed by Nelder and Mead (1965). For these figures, we adjust the plots so that the horizontal and vertical scales are the same within each column (but not within each row).

The first row considers the CG method. Looking only at the right histogram, we see that at least half of the initial values result in parameter estimates that give log likelihood values more than 100 ($= .1 \times 1000$) less than the log likelihood value at the model parameters. The second row reports on the performance of the Nelder-Mead algorithm. It actually performs fairly well, but the right histogram shows that there are roughly 300 starting values at which the log likelihood is more than 31.6 ($= 10^{-1.5} \times 1000$) less than the log likelihood value at the model parameters.

We also implemented a simulated annealing algorithm, specifically a variant given by Belisle (1992), which we refer to in the paper as SANN. Due to its less common usage in political science, we present those results in the appendix. (See Figure 12.) Briefly, the SANN algorithm performs very well, with all but 41 of the 1000 initial values leading to estimates for which the log likelihood values are greater than the log likelihood values at the model parameters.

In addition, we looked at a few alternative situations. For one, we considered the case of continuous values of the data for the \mathbf{z} variables, rather than the discrete values used in the simulated data set discussed earlier. (We still take the components of \mathbf{z} to be nonnegative.) Running search algorithms on a simulated data set with continuous z variables led to problems like those found above, namely estimated parameters far from the model parameters, and values of the log likelihood function at the estimated parameters far from its values at the model parameters.⁹

We also looked at some simulations in which the model parameters γ were a bit smaller in absolute value. Recall that the model parameter γ_0 in the data set above is $(3.43, -4.07)$.

⁹We note that our intuitive explanation given in Section 3.1 indicates that these problems would arise for both discrete and continuous \mathbf{z} . In fact, that reasoning might lead us to believe the problems could be even worse for continuous \mathbf{z} , given that we would expect far fewer occurrences of $\mathbf{z} = \mathbf{0}$ in the continuous case.

The reason for doing so is that one could be concerned that this value would lead to a large variance ($\exp(3.43)$) in the case that $\mathbf{z} = (1, 0)$ and a small variance ($\exp(-4.07)$) in the case that $\mathbf{z} = (0, 1)$; on the other hand, in the case that $\mathbf{z} = (1, 1)$, the variance ($\exp(3.43 - 4.07)$) would not be as large or small. There were similar problems with the search algorithms for a data set with $\gamma_0 = (-0.6, 0.84, -0.69, -0.15, -0.16, 0.42)$, although the problems did occur with less frequency (as a fraction of the number of choices of initial values).¹⁰ We also note that below we find similar problems with a real data set used by Alvarez and Brehm, which has a parameter γ whose components are each smaller than 1 in absolute value.¹¹

2.2 Simulated Data: Probit Model

We now contrast the results for the above search algorithms for the heteroskedastic probit model with their performance for a probit model with the same number of parameters.

Thus, we consider a probit model with $\dim(\beta) = 5$, so that the dimension of the parameter vector is the same in both cases. We again generate a pseudorandom data set with 1000 observations \mathbf{x}_i , with components equal to 0 or 1 (and with the first component always equal to 1). Then we generate a choice of β_0 by selecting the second through fifth components of β_0 from $[-5, 5]$, but we choose the first component so that the mean of $\mathbf{x}_i' \beta_0$ over all observations is 0. In the probit setting, we do not make any requirements about the fraction of crossover. We again generate 1000 choices for initial parameters in the box $[-5, 5]^5$ and start each search technique at all 1000 of these choices.

Results for the BFGS search algorithm are presented in Figure 3. All 1000 initial values lead to estimates at which the log likelihood is greater than at the model parameters. Moreover, all of the estimates are in fact within .233 of the model parameters. These graphs can be compared directly with Figure 1.

For the CG and Nelder-Mead search algorithms, the results are similar. We present the results in Figure 4, which can be compared directly with Figure 2. For the CG algorithm, as

¹⁰We note that, in very limited testing, we did not find any problems with local search for simulated data sets with $\dim(\gamma_0) = 2$.

¹¹These are not exactly a model parameters in this case, but is analogous— see below.

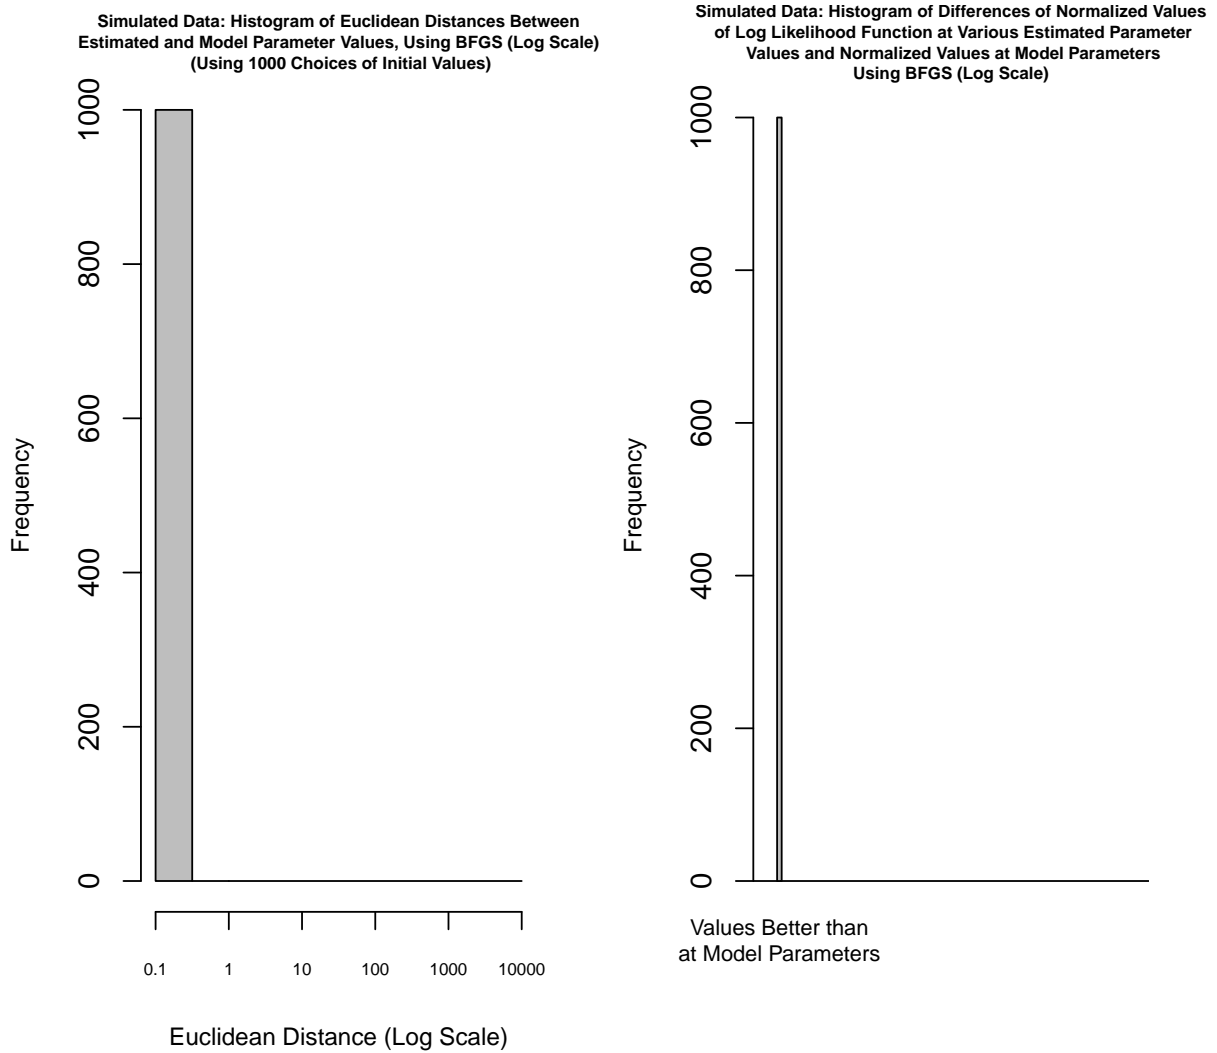


Figure 3: Performance of the BFGS Search Algorithm on a Simulated Data Set for the Probit Model, for 1000 Random Choices of Initial Values

with the BFGS method, all 1000 initial values lead to estimates at which the log likelihood is greater than at the model parameters. The Nelder-Mead algorithm leads to some results that are not as good; however, all but 19 initial values lead to estimates at which the log likelihood is better than at the model parameters.¹² The performance of these search algorithms when applied to the probit model is far better than the analogous performance

¹²The results for the SANN algorithm are presented in Figure 13, in the Appendix.

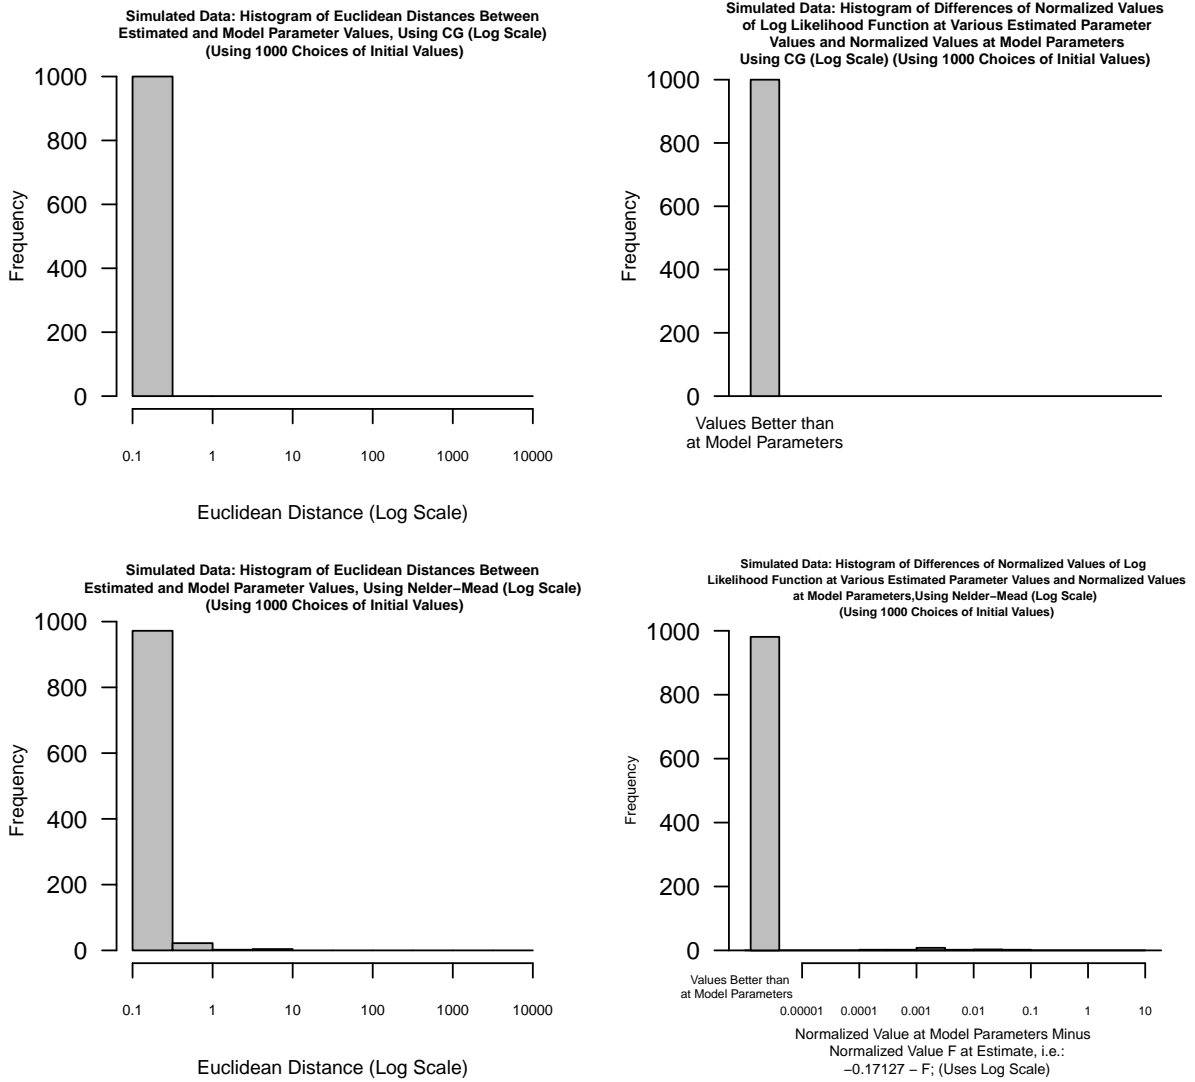


Figure 4: Performance of the CG and Nelder-Mead Search Algorithms on a Simulated Data Set for the Probit Model, for 1000 Random Choices of Initial Values

for the heteroskedastic probit model.

2.3 Alvarez-Brehm Data: Heteroskedastic Probit Model

We now consider the behavior of search techniques when employed on a data set used by Alvarez and Brehm (1995) in their study of attitudes toward abortion. We use the same specification as Alvarez and Brehm, namely a heteroskedastic probit model using a constant

and seven variables in what they call the “choice model” (informally, the part depending on β , so that the dimension of β is eight), and six variables in the “variance model” (loosely, the part depending on γ , so that the dimension of γ is six). Alvarez and Brehm actually use these explanatory variables in seven models, one for each of seven outcome variables. Here we focus on one of these outcomes: whether or not participants think that abortion should be legal if a woman wants it for any reason.¹³ In the data, we remove any observations for which one of the explanatory or outcome variables is missing; in so doing, we match Alvarez and Brehm in using 1295 observations.

As above, we are interested in how close the estimates are to optimal in some sense. For operational purposes, we take the optimal value to be the estimates found by Alvarez and Brehm, which we refer to as the A-B parameter values. Denote this set of parameter values by $(\beta_{AB}, \gamma_{AB})$ ¹⁴. We now present histograms of the two measures used above to assess the performance of the search methods. Here, for the normalized measure, we will divide by $n_{AB} = 1295$, the number of observations in the Alvarez-Brehm data. We note that $\frac{1}{n_{AB}}F((\beta_{AB}, \gamma_{AB})) = -0.59383$. We begin with the BFGS algorithm. The two histograms are displayed in Figure 5. Both plots use a log scale on the horizontal axis, to more clearly present key phenomena; as above, this log scale only applies to the right portion of the right plot.¹⁵

Recall that, viewing the Alvarez-Brehm parameters as correct in some sense, we expect the log likelihood to be in general greater (i.e. negative and having less magnitude) at these parameters than at our estimates. However, for roughly 300 choices of starting values, the values of the log likelihood at the estimates found using the BFGS algorithm are actually greater than the value at the Alvarez-Brehm parameters. This is most likely due to the fact that the parameters we use as the Alvarez-Brehm parameters are from their paper, and

¹³See Alvarez and Brehm (1995) for more details. See Table 1 of their paper for results for the seven models.

¹⁴These estimates can be found in the final column of Table 1 of their paper: they are $\beta_{AB} = (-.07, -.15, -.13, .05, -.22, -.79, .12, .51)$ and $\gamma_{AB} = (-.22, -.48, .22, -.30, .68, .63)$.

¹⁵Observe that, as with the graphs for the simulated data, the scales for the horizontal and vertical axes differ by plot.




Figure 5: Performance of the BFGS Search Algorithm on the Alvarez-Brehm Data (Using the Heteroskedastic Probit Model), for 1000 Random Choices of Initial Values

reported only to two digits.¹⁶ Thus we are probably essentially getting the same parameter estimates as they are, as casual inspection of a few of the better estimates indicates. As above, we represent this in the graph by plotting a rectangle on the left side of the plot with height equal to the number of these cases and labeled as “Values Better than at A-B

¹⁶Note that this is unlike the situation for the simulated data discussed above, as here we would expect that Alvarez and Brehm found the estimates that maximize the log likelihood function, rather than finding the estimates that maximize the analogous expectation, which would typically not be maximized at the model parameters used to generate a simulated data set.

Parameters.”¹⁷

When starting with random initial values, the parameter estimates only appear to be close to the A-B parameters for about 30% of the 1000 runs. In fact, the Euclidean distance between the parameter estimates and the A-B parameters is at least 10 for 683 of the 1000 runs. Moreover, the *normalized* value of the log likelihood function at the estimated parameter values is at least .063 less than the *normalized* value of the log likelihood function at the A-B parameters for more than 600 of the runs; thus the (unnormalized) value of the log likelihood function is at least 81.5 less than the (unnormalized) value at the A-B parameters for these runs.¹⁸ Thus by these two measures the BFGS algorithm appears to perform well for only a fraction of the randomly chosen starting values.

Figure 6 presents similar plots, again for the Alvarez-Brehm data set and starting at each of the same 1000 choices of initial values, but now we look at the CG and Nelder-Mead algorithms. As above, for these figures, we adjust the plots so that the horizontal and vertical scales are the same within each column (but not within rows).

From the right plot in the top row, one can see that at least roughly 600 of the results from the CG method appear to be fairly far from optimal, as judged using the value of the log likelihood function at these estimates. There are almost about 100 choices of initial value however that yield estimates better than the A-B parameter values by this measure, and a few that are only slightly worse (about 10^{-9} or 10^{-6} larger). From the left plot in the bottom row, all but about 100 of the estimated parameters from the Nelder-Mead method have Euclidean distance at least 1 from the A-B parameter values; this may not necessarily be indicative of poor estimates, but from the right plot we can see that at least about 100 estimates are not close, in that the normalized values are at least 0.1 larger than the normalized value at the A-B parameters. (So the log likelihood function is at least 12.95 larger.) In the Appendix, we look at the performance of the SANN search algorithm on the A-B data. (See Figure 14.) The results are neither clearly better nor worse than those for

¹⁷Again, its location along the horizontal axis is of course not meaningful.

¹⁸The value .063 equals $10^{-1.2}$, which is the left endpoint of the rightmost histogram cell in the plot on the right in Figure 5; the value 81.5 is just less than the product $1295 \times .063$.

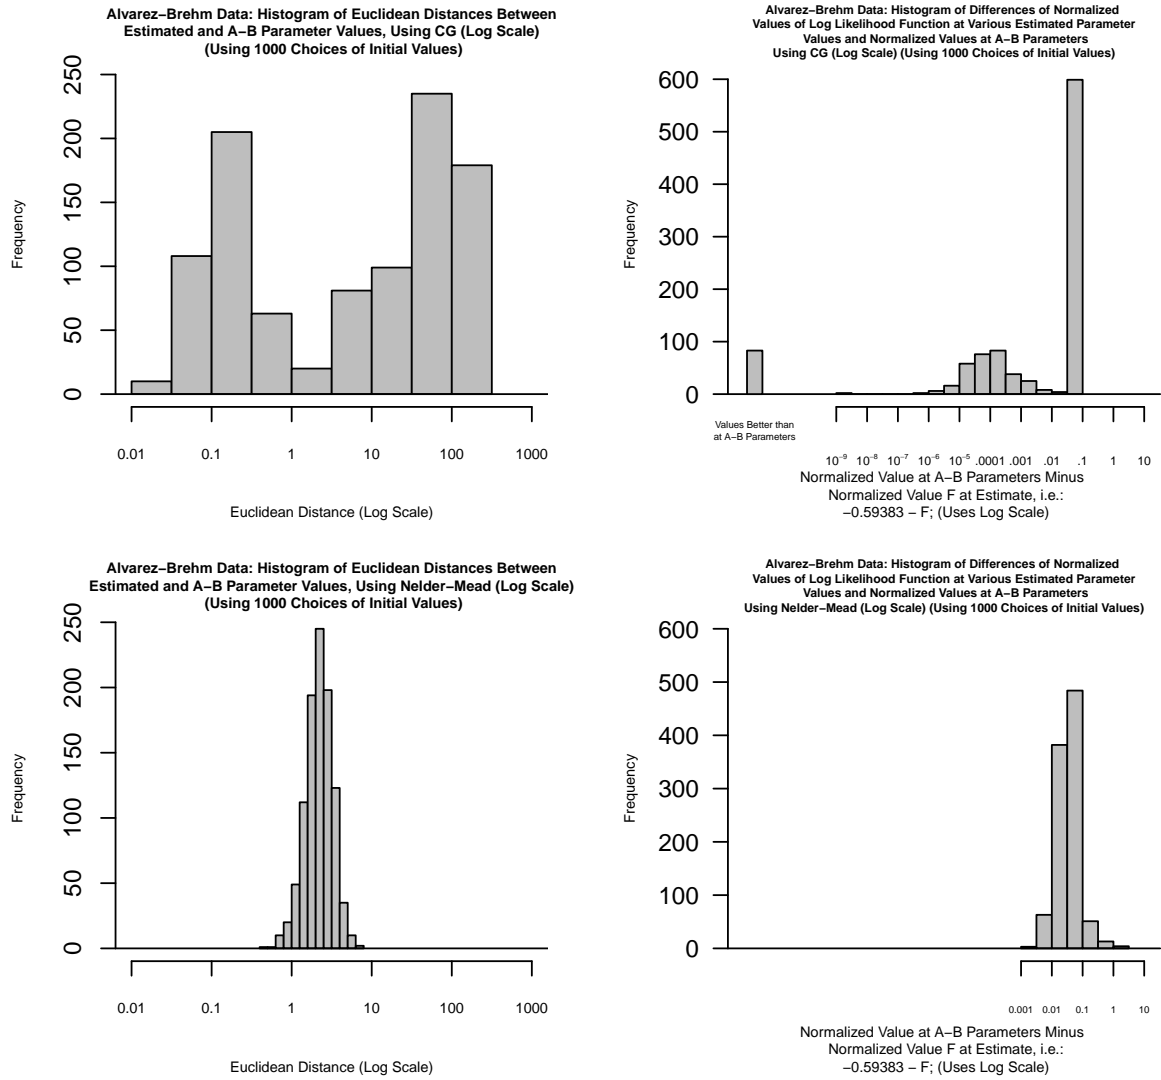


Figure 6: Performance of the CG and Nelder-Mead Search Algorithms on the Alvarez-Brehm Data (Using the Heteroskedastic Probit Model), for 1000 Random Choices of Initial Values

the other search techniques.

3 Observations on the Shape of the Log Likelihood Function for the Heteroskedastic Probit

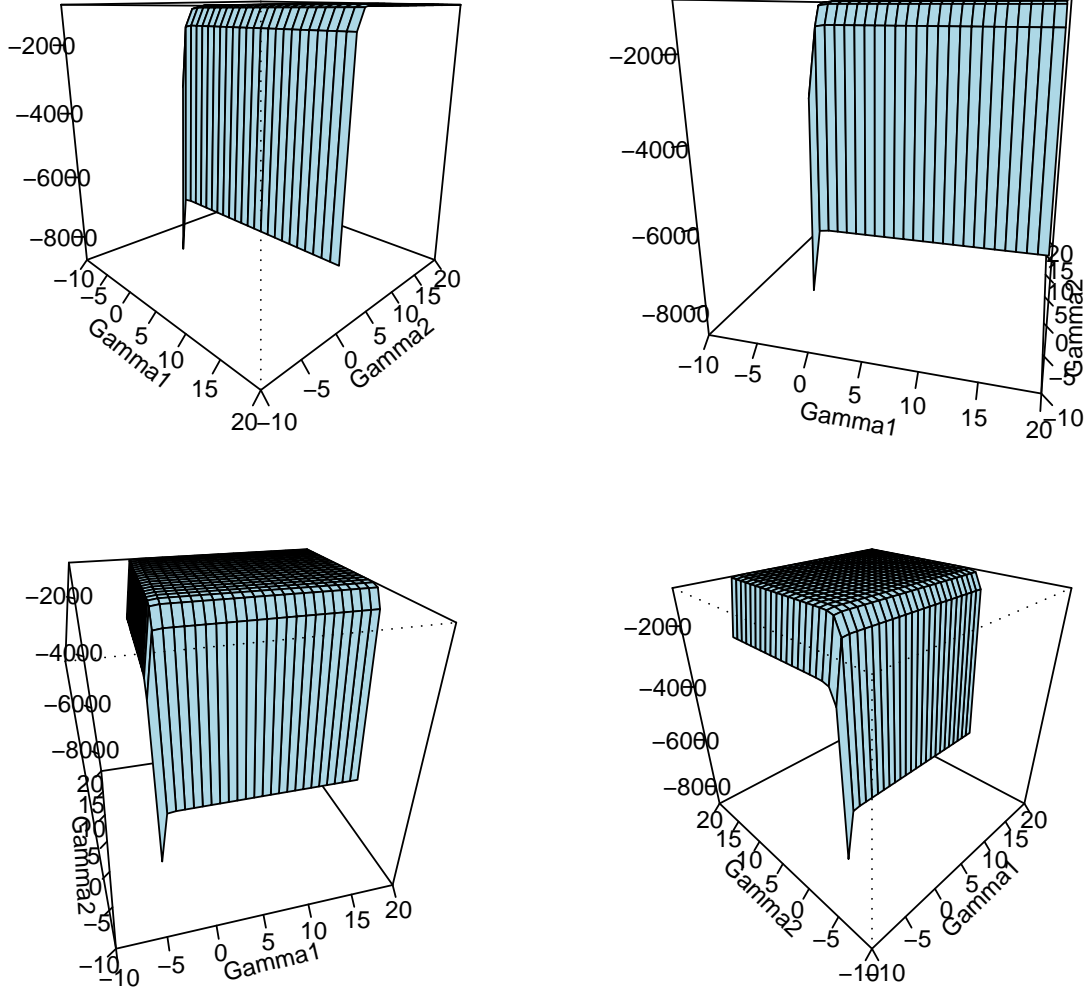
Next, we explore why optimization algorithms might perform poorly when applied to the heteroskedastic probit likelihood. We start our exploration with a profile plot of the log likelihood function for the heteroskedastic probit using the the Alvarez-Brehm data. We first fix a parameter estimate arising from the BFGS algorithm for a particular initial value, which we have chosen to have a resulting parameter estimate at which the value of the log likelihood is not close to its value at the Alvarez-Brehm estimates. We display a graph that demonstrates how the log likelihood varies as the first two components γ_1 and γ_2 of γ vary.¹⁹ The value of β and the other components of γ are fixed at the estimated values. Figure 7 presents this plot, from four different vantage points. In order to make the graph readable, we restrict the range of output values plotted to be those values at least $-10,000$. At the values of (γ_1, γ_2) for which no output values are plotted, the log likelihood function is actually much less than $-10,000$.

We see a distinct shape in this graph: a plateau, which ranges roughly over positive values of (γ_1, γ_2) . This gives a sense of the plateau but investigation of the graph restricted to a smaller range of output values shows that it has a more complex structure. For example, there is a ridge for γ_2 fixed at a certain value a little less than zero, which declines in γ_1 . If this plateau were in some sense flat enough, an optimization algorithm that is essentially local in nature may not be able find the actual maximum of the function if it starts at an initial value somewhere on the plateau. Additionally, it may be led to this plateau from other regions and then subsequently remain there.

¹⁹Throughout Section 3, for ease of notation, we often use β and γ where $\hat{\beta}$ and $\hat{\gamma}$ might perhaps be more appropriate.

Figure 7:

Graphs of (γ_1, γ_2) -Slice of Log Likelihood Function from Different Perspectives,
for a Restricted Range of Output Values, for the Alvarez-Brehm Data Set at Estimated Parameters



Notes: Only output values greater than or equal to -10000 are graphed. The value of the log likelihood at the A-B Parameters is: -769.01 . The estimated parameters are:
 $\beta = (-65.34, -55.53, -40.91, -7.87, 2.41, -34.08, 78.8, 62.68)$.
 $\gamma = (50.44, 38.75, 108.81, 77.17, 23.58, -22.99)$.
 The initial values are:
 $\beta_{init} = (-0.06, -0.96, -0.21, 0.99, 0.8, -0.27, 0.17, 0.83)$.
 $\gamma_{init} = (-0.82, -0.34, 0.05, -0.26, 0.73, 0.93)$.

3.1 The Heteroskedastic Probit Log Likelihood Function

Next we discuss why we might expect a plateau region of the sort seen in Figure 7. Consider the form of the log likelihood function for the heteroskedastic probit, given in Equation (2).

In this section, we focus on the case in which \mathbf{z} is discrete and, for all i with $1 \leq i \leq n$, all of the components of \mathbf{z}_i are nonnegative. For convenience, we denote this by the shorthand

$$Z \geq 0,$$

where we are thinking of Z as above, i.e. as the $n \times k_2$ matrix with i^{th} row given by \mathbf{z}'_i . In the Alvarez-Brehm data used above, this holds, and it is also the case in the simulated data in Section 2. We will return to consideration of more general \mathbf{z} later. We will also assume in this section that there are some values i for which $\mathbf{z}_i = \mathbf{0}$.

Then, except in the case that $\mathbf{z}_i = \mathbf{0}$, for values of $\boldsymbol{\gamma}$ with components that are positive and sufficiently large, we can ensure that $\mathbf{z}'_i \boldsymbol{\gamma}$ is large and positive. Of course for such i , we also have that $\exp(\mathbf{z}'_i \boldsymbol{\gamma})$ is large, whence the term $\frac{\mathbf{x}'_i \boldsymbol{\beta}}{\exp(\mathbf{z}'_i \boldsymbol{\gamma})}$ will be close to 0. Thus, under these conditions we have

$$\begin{aligned}
\ell(\boldsymbol{\beta}, \boldsymbol{\gamma} \mid \mathbf{y}, X, Z) &\approx \sum_{1 \leq i \leq n: \mathbf{z}_i = \mathbf{0}} (y_i \ln \Phi(\mathbf{x}'_i \boldsymbol{\beta}) + (1 - y_i) \ln(1 - \Phi(\mathbf{x}'_i \boldsymbol{\beta}))) \\
&\quad + \sum_{1 \leq i \leq n: \mathbf{z}_i \neq \mathbf{0}} (y_i \ln \Phi(0) + (1 - y_i) \ln(1 - \Phi(0))) \\
&\approx \sum_{1 \leq i \leq n: \mathbf{z}_i = \mathbf{0}} (y_i \ln \Phi(\mathbf{x}'_i \boldsymbol{\beta}) + (1 - y_i) \ln(1 - \Phi(\mathbf{x}'_i \boldsymbol{\beta}))) \\
&\quad - (\ln 2) \times \left| \{ \mathbf{z}_i : \mathbf{z}_i \neq \mathbf{0}, \quad 1 \leq i \leq n \} \right|. \tag{4}
\end{aligned}$$

From the form of the expression in (4), we can see why there is a plateau in the graphs in Figure 7. Consider a choice $\boldsymbol{\gamma}_1 = (\gamma_{11}, \gamma_{12}, \dots, \gamma_{1k_2})$ with large positive components. Throughout the remainder of the paper, we at times abuse language and refer to such vectors $\boldsymbol{\gamma}$ informally as simply “large and positive.” Then, for $\boldsymbol{\gamma} = (\gamma_{11}, \gamma_{12}, \dots, \gamma_{1k_2})$ near $\boldsymbol{\gamma}_1$, the denominator $\exp(\mathbf{z}'_i \boldsymbol{\gamma})$ will still be large and positive, so that $\mathbf{x}'_i \boldsymbol{\beta} / \exp(\mathbf{z}'_i \boldsymbol{\gamma})$ will be small and relatively unaffected by small changes in $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$.

To see this in more detail, assume that there is a choice of parameters $(\boldsymbol{\beta}_1, \boldsymbol{\gamma}_1)$ and some large value of $L > 0$ so that $\boldsymbol{\gamma}_1$ satisfies

$$\min_{1 \leq j \leq k_2} \gamma_{1j} \geq L. \tag{5}$$

Recall that we have assumed that the \mathbf{z} ’s are discrete and have nonnegative components.

Now assume for convenience that in fact the smallest positive value of the components of the \mathbf{z} 's is 1.

Consider a new choice of parameters $(\boldsymbol{\beta}_2, \boldsymbol{\gamma}_2)$ with $\boldsymbol{\gamma}_2 = (\gamma_{21}, \gamma_{22}, \dots, \gamma_{2k_2})$ satisfying, say,

$$|\gamma_{2j} - \gamma_{1j}| \leq 1 \quad \text{for } 1 \leq j \leq k_2. \quad (6)$$

Then (5) and (6) together imply that

$$\min_{1 \leq j \leq k_2} \gamma_{2j} \geq L - 1.$$

Also assume that $H > 0$ satisfies

$$|\mathbf{x}'_i \boldsymbol{\beta}_1| \leq H \quad \text{and} \quad |\mathbf{x}'_i \boldsymbol{\beta}_2| \leq H \quad \text{for } 1 \leq i \leq n.$$

Now we consider terms of the log likelihood function for which \mathbf{z}_i is not zero. Fix any i_1 with $1 \leq i_1 \leq n$ satisfying $\mathbf{z}_{i_1} \neq \mathbf{0}$. As we have assumed that \mathbf{z} is discrete with a minimum nonzero value of 1, it follows that $\mathbf{z}'_{i_1} \boldsymbol{\gamma}_2 \geq L - 1$, and thus $\exp(\mathbf{z}'_{i_1} \boldsymbol{\gamma}_2) \geq \exp(L - 1)$ holds. Therefore one has

$$\left| \frac{\mathbf{x}'_{i_1} \boldsymbol{\beta}_j}{\exp(\mathbf{z}'_{i_1} \boldsymbol{\gamma}_j)} \right| \leq H e^{1-L} \quad \text{for } j = 1 \text{ and } j = 2.$$

So certainly for L large relative to H ,

$$\frac{\mathbf{x}'_{i_1} \boldsymbol{\beta}_1}{\exp(\mathbf{z}'_{i_1} \boldsymbol{\gamma}_1)} \quad \text{and} \quad \frac{\mathbf{x}'_{i_1} \boldsymbol{\beta}_2}{\exp(\mathbf{z}'_{i_1} \boldsymbol{\gamma}_2)}$$

will both be small in magnitude, and thus both Φ and $(1 - \Phi)$ evaluated at these expressions will be very close to $1/2$ (as $\Phi(0) = 1/2$). So the values of the i_1^{th} summand of the log likelihood at the two parameter choices $(\boldsymbol{\beta}_1, \boldsymbol{\gamma}_1)$ and $(\boldsymbol{\beta}_2, \boldsymbol{\gamma}_2)$ will clearly be very close to each other; that is, the i_1^{th} term of the log likelihood function will not change substantially for small perturbations of $\boldsymbol{\gamma}$ with large, positive components. Of course, for i such that $\mathbf{z}_i = \mathbf{0}$, the i^{th} term of the log likelihood could change in a substantial way if $\boldsymbol{\gamma}$ is on the

plateau. But in certain situations, these terms may be a small portion of the total number of terms, and thus may not have a strong effect on the value of the sum as a whole.

Now consider an optimization algorithm used to maximize the log likelihood. Suppose the algorithm reaches a point (β, γ) on the plateau, i.e. with γ having large and positive components. As seen above, any small changes in γ , regardless of the change in β , will not substantially affect the terms indexed by i for which $\mathbf{z}_i \neq \mathbf{0}$. We propose that the effect of changes in the remaining terms might not be enough to guide the algorithm to the best value; any improvements in the terms indexed by those i for which $\mathbf{z}_i = \mathbf{0}$ that might be caused by veering off of the plateau could perhaps be outweighed by increased penalties in the far greater number of terms indexed by i such that $\mathbf{z}_i \neq \mathbf{0}$.

Before turning to another topic, we make another observation. As discussed above, if γ has large and positive components, then for i with $\mathbf{z}_i \neq \mathbf{0}$, the expression $\frac{\mathbf{x}'_i \beta}{\exp(\mathbf{z}'_i \gamma)}$ will be close to 0. Thus $\Phi\left(\frac{\mathbf{x}'_i \beta}{\exp(\mathbf{z}'_i \gamma)}\right)$ and $\left(1 - \Phi\left(\frac{\mathbf{x}'_i \beta}{\exp(\mathbf{z}'_i \gamma)}\right)\right)$ will be close to 1/2. Hence the terms in the log likelihood indexed by i for which \mathbf{z}_i is nonzero will be close to $-\ln 2$. Thus, we expect to see values of the log likelihood function close to $-n \ln 2$ at the estimates on the plateau, especially if there are very few terms with $\mathbf{z}_i = \mathbf{0}$. Thus, in some sense, we can also think of $-n \ln 2$ as a sort of benchmark for a value that the optimization algorithm should be able to easily achieve. Observe that $-\ln 2 \approx -0.693$. One can compare the values obtained in the simulation to this value. For example, consider the right plot in Figure 5. The normalized value of the log likelihood function at the A-B parameters is $-.59383$, as can be read off from the horizontal axis of the plot. There are more than 600 initial values from which the BFGS search method results in estimates at which the normalized value of the log likelihood is almost .1 less than the value at the A-B parameters, which gives a normalized value on the order of $-\ln 2$. As we expect for the reasons outlined here, it was in fact common in our various computations for many of these normalized values to be fairly close to $-\ln 2$ at plateau solutions.

Considering this benchmark gives another way to think intuitively about the benefit to an optimizer of choosing a large positive value for γ . The function $\ln(\Phi(x))$ decreases rapidly

for x negative and decreasing. (For example, $\ln(\Phi(-10)) \approx -53$.) So, in some sense, the penalty for missing by a lot for one observation can be very large, whereas the gain from improving the other observations may not be that large. So it may be beneficial to simply make every term fit at least somewhat decently, and a large positive value for γ allows us to at least limit these penalties to around -0.693 each. This is obviously not a rigorous line of reasoning, but it gives a heuristic for why search techniques might often end with a plateau solution.

3.1.1 Graphical Interpretation: The Basic Probit Model

Now we give a graphical interpretation that sheds light on the underlying mechanics of how optimization algorithms perform, when used to maximize the log likelihood function associated with the heteroskedastic probit model. To develop the interpretation, we first consider a probit model with two-dimensional parameter β , so that $\Pr(Y = 1 \mid \mathbf{x}) = \mathbf{x}'\beta$. For a fixed choice $\beta \in \mathbb{R}^2$, Figure 8 depicts several observations $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_5 \in \mathbb{R}^2$ and the hyperplane $\mathbf{x}'\beta = 0$, which of course in \mathbb{R}^2 is simply a line. Around each observation \mathbf{x}_i , we draw a dashed line segment perpendicular to the line $\mathbf{x}'\beta = 0$, extending in either direction from the point \mathbf{x}_i ; the endpoints are those points \mathbf{v}_{i1} and \mathbf{v}_{i2} on this perpendicular line that satisfy

$$\mathbf{v}_{i1}'\beta = \mathbf{x}_i'\beta - 1 \quad \text{and} \quad \mathbf{v}_{i2}'\beta = \mathbf{x}_i'\beta + 1.$$

This dashed line represents adding an error term ϵ_i to the linear combination $\mathbf{x}_i'\beta$. The values of $\mathbf{v}'\beta$ for vectors \mathbf{v} correspond to values of the latent variable $y_i^* = \mathbf{x}_i'\beta + \epsilon_i$ for $|\epsilon_i| \leq 1$.

The region $\{\mathbf{x} : \mathbf{x}'\beta > 0\}$ is, in this example, the upper right half plane in Figure 8, because the value of β we use is $(1.5, 1)$. This region roughly corresponds to the region with those \mathbf{x}_i for which the associated y_i equals 1, while the lower left half plane is the region $\{\mathbf{x} : \mathbf{x}'\beta < 0\}$, which corresponds roughly to the region with $y_i = 0$; note that we say “roughly” because one could have what we have termed crossover.²⁰ Crossover occurs

²⁰We ignore points \mathbf{x}_i satisfying $\mathbf{x}_i'\beta = 0$ here as its consideration would only complicate the discussion,




Figure 8: A Graphical Representation of the Probit Model

if either (i) $\mathbf{x}'_i\boldsymbol{\beta} > 0$ holds yet $y_i = 0$, or (ii) $\mathbf{x}'_i\boldsymbol{\beta} < 0$ holds yet $y_i = 1$. For an example, consider \mathbf{x}_2 in Figure 8. One in fact has $\mathbf{x}'_2\boldsymbol{\beta} = -0.5$. If we assume that $\epsilon_2 = 0.75$, then the latent variable $y_2^* = \mathbf{x}'_2\boldsymbol{\beta} + \epsilon_2 = 0.25$, so $y_2 = 1$, and thus for this value of ϵ_2 , observation 2 exhibits crossover.

Now consider, in terms of Figure 8, what choices might be made by an optimization algorithm used to maximize the value of the log likelihood for the probit model. Despite which we give only for the sake of intuitive understanding, and moreover would not add any material value.

the fact that generally an algorithm is used to maximize the log likelihood, for exposition purposes, we will generally speak of an “optimizer” as if an individual were attempting to maximize the log likelihood. So consider the task facing this optimizer. One key goal, given data $\{(y_i, \mathbf{x}_i)\}_{1 \leq i \leq n}$, is to choose $\boldsymbol{\beta}$ so that most \mathbf{x}_i ’s are on the correct side of $\mathbf{x}_i' \boldsymbol{\beta} = 0$, so to speak; i.e., one would choose \mathbf{x}_i to be in the half-plane $\mathbf{x}_i' \boldsymbol{\beta} > 0$ if $y_i = 1$ and in the half-plane $\mathbf{x}_i' \boldsymbol{\beta} < 0$ if $y_i = 0$. In other words, an optimizer would aim to minimize crossover. Of course, an optimizer need not do so exactly.²¹ We can think of each summand of the log likelihood function as assigning a penalty if there is crossover; of course, there is also a summand that is added even if there is no crossover for the observation, but all of these summands are negative and we think of most of them as having small magnitude, so that intuitively, an optimizer might be mainly concerned with the summands for which there is in fact crossover. These penalties for crossover are determined as follows. First, consider what we have just noted: this penalty is always negative. For the points \mathbf{x}_i that are far from the line $\mathbf{x}_i' \boldsymbol{\beta} = 0$, the penalty for crossover in the log likelihood function is much larger (in magnitude) than the penalty for points \mathbf{x}_i near the line. For example, the i^{th} term of the log likelihood function for i with $y_i = 1$ is $\ln \Phi(\mathbf{x}_i' \boldsymbol{\beta})$, which is of course much more negative, the more negative is the value of $\mathbf{x}_i' \boldsymbol{\beta}$. So, for example, if $y_1 = y_2 = 1$ were to hold, then given their positions as shown in Figure 8, it is more important to put \mathbf{x}_1 in the correct half plane than \mathbf{x}_2 , as the penalty $\ln \Phi(\mathbf{x}_1' \boldsymbol{\beta})$ from not doing so is much larger (in magnitude) than the corresponding penalty $\ln \Phi(\mathbf{x}_2' \boldsymbol{\beta})$.

3.1.2 Graphical Interpretation: The Heteroskedastic Probit Model

Now we consider the heteroskedastic probit model, and, for simplicity, we examine the case in which $\dim(\mathbf{x}) = 2$ and $\dim(\mathbf{z}) = 2$, i.e. for $\boldsymbol{\beta}, \boldsymbol{\gamma} \in \mathbb{R}^2$. The picture here includes a few modifications to Figure 8. Again, the objective of an optimizer involves choosing $\boldsymbol{\beta}$ to minimize crossover but now with an additional tool, namely the optimizer can choose $\boldsymbol{\gamma}$ so as to affect the penalties associated with crossover. Figure 9 presents the graphs. The top

²¹And indeed it is probably extremely unlikely with real-world data to be able to do so, and if so, there would be the problem of complete separation.




Figure 9: A Graphical Representation of the Heteroskedastic Probit Model

graph is similar to the plot in Figure 8. Again, the line $\mathbf{x}'\underline{\beta} = 0$ separates the observations $\mathbf{x}_1, \dots, \mathbf{x}_5$ into two half-planes. Also, we again draw dashed line segments around each \mathbf{x}_i ,

in the direction perpendicular to the line $\mathbf{x}'\boldsymbol{\beta} = 0$. But now these line segments extend for lengths that vary for each observation; this length, for the point \mathbf{x}_i , varies with $\exp(\mathbf{z}'_i\boldsymbol{\gamma})$; much as in the probit graph, the endpoints of the i^{th} line segment are the points \mathbf{v}_{i1} and \mathbf{v}_{i2} satisfying

$$\mathbf{v}'_{i1}\boldsymbol{\beta} = \mathbf{x}'_i\boldsymbol{\beta} - \exp(\mathbf{z}'_i\boldsymbol{\gamma}) \quad \text{and} \quad \mathbf{v}'_{i2}\boldsymbol{\beta} = \mathbf{x}'_i\boldsymbol{\beta} + \exp(\mathbf{z}'_i\boldsymbol{\gamma}). \quad (7)$$

Recall from (1) that $\exp(\mathbf{z}'_i\boldsymbol{\gamma})$ is the standard deviation²² of the error term ϵ_i that forms part of the definition of the latent variable y_i^* . For those observations i with $\exp(\mathbf{z}'_i\boldsymbol{\gamma})$ large, speaking very loosely, the model assigns a larger probability of crossover (speaking very loosely again), as the error term is more likely to be larger, as represented in the picture by a longer dashed line segment.

Consider our earlier framework of thinking of each summand of the log likelihood function as giving a penalty for crossover (and contributing a small negative number to the sum if not). Recall from (2) that this penalty is always negative. Suppose that there is crossover for observations 1 and 2; from the graph, we can see that these two points are in the region $\{\mathbf{x} : \mathbf{x}'\boldsymbol{\beta} < 0\}$, so in other words, given that we are assuming there is crossover, we have $y_1 = y_2 = 1$.²³ From (2), the penalty for crossover, as $y_1 = y_2 = 1$, is

$$\ln \left(\Phi \left(\frac{\mathbf{x}'_i\boldsymbol{\beta}}{\exp(\mathbf{z}'_i\boldsymbol{\gamma})} \right) \right) \quad \text{for } i = 1, 2,$$

which is, speaking very loosely, smaller (in magnitude), the larger is $\exp(\mathbf{z}'_i\boldsymbol{\gamma})$ (as $\mathbf{x}'_i\boldsymbol{\beta} < 0$ for $i = 1, 2$). Of course, the size of the penalty also depends on the magnitude of $\mathbf{x}'_i\boldsymbol{\beta}$, and more precisely on the ratio $\frac{\mathbf{x}'_i\boldsymbol{\beta}}{\exp(\mathbf{z}'_i\boldsymbol{\gamma})}$. It is not immediately clear from the graph, but for our purpose here of providing exposition, one can perhaps see that it makes some intuitive sense that this penalty would be smaller (in magnitude) for observation 1 than for observation 2.²⁴

²²The decision to use the standard deviation is somewhat arbitrary; we are merely showing one pictorial representation of the main concepts.

²³As with the probit graph, one needs to know the value of $\boldsymbol{\beta}$ to determine which side of the line is the region $\{\mathbf{x} : \mathbf{x}'\boldsymbol{\beta} < 0\}$ and which side the region $\{\mathbf{x} : \mathbf{x}'\boldsymbol{\beta} > 0\}$. Here we have in fact chosen $\boldsymbol{\beta} = (-1.2, 1)$. Obviously had we chosen $(1.2, -1)$, or some positive scalar multiple of it, the regions would be reversed.

²⁴To see this precisely, first recall that we have $\mathbf{x}'_1\boldsymbol{\beta} < 0$ and $\mathbf{x}'_2\boldsymbol{\beta} < 0$. The fact that the dashed line for observation 1 crosses the line $\mathbf{x}'\boldsymbol{\beta} = 0$ shows, from (7), that $\exp(\mathbf{z}'_1\boldsymbol{\gamma}) > |\mathbf{x}'_1\boldsymbol{\beta}|$, whence $\mathbf{x}'_1\boldsymbol{\beta} / \exp(\mathbf{z}'_1\boldsymbol{\gamma}) > -1$

Thus a maximizer gets penalized much less for choosing β so that \mathbf{x}_1 is on the wrong side of the dividing line than for choosing β so that \mathbf{x}_2 is on the wrong side.

Now consider the bottom graph in Figure 9, which depicts the observations \mathbf{z}_i ($1 \leq i \leq 5$). Also we see the hyperplane $\mathbf{z}'\gamma = 0$ (here a line, as we have assumed above that γ is two-dimensional). Here the upper right half represents the region $\{\mathbf{z} : \mathbf{z}'\gamma > 0\}$, and the lower left half plane the region $\{\mathbf{z} : \mathbf{z}'\gamma < 0\}$.²⁵ Note that there is no special significance of the line $\mathbf{z}'\gamma = 0$ analogous to the significance of the line $\mathbf{x}'\beta = 0$, but we include it as an indicator of the value of $\mathbf{z}'\gamma$. Observe that those observations i with $\mathbf{z}'_i\gamma$ positive and large will have $\exp(\mathbf{z}'_i\gamma)$ large and thus have a longer associated dashed line segment in the top graph of Figure 9.

An optimizer, then, attempts to simultaneously choose two lines (or, more generally, two hyperplanes): one to split the \mathbf{x}_i 's into two regions to try to give observations a positive value for $\mathbf{x}'_i\beta$ if $y_i = 1$ and a negative value if $y_i = 0$, and the second to arrange the \mathbf{z}_i 's so that (loosely speaking) the smaller is $\mathbf{z}'_i\gamma$, the less likely observation i is to exhibit crossover.

Now observe what happens for an optimizer who chooses a plateau solution. In this case, we return to the earlier assumption that $Z \geq 0$ (which is of course not the case in the bottom graph in Figure 9), and that there are some observations i with $\mathbf{z}_i = \mathbf{0}$. For this

holds. On the other hand, the dashed line for observation 2 does not cross the line $\mathbf{x}'\beta = 0$. Thus $\exp(\mathbf{z}'_2\gamma) < |\mathbf{x}'_2\beta|$, whence $\mathbf{x}'_2\beta / \exp(\mathbf{z}'_2\gamma) < -1$ holds. Therefore we have $\mathbf{x}'_2\beta / \exp(\mathbf{z}'_2\gamma) < \mathbf{x}'_1\beta / \exp(\mathbf{z}'_1\gamma)$ and thus $\ln(\Phi(\mathbf{x}'_2\beta / \exp(\mathbf{z}'_2\gamma))) < \ln(\Phi(\mathbf{x}'_1\beta / \exp(\mathbf{z}'_1\gamma)))$. So the penalty for observation 2 is more negative than that for observation 1, i.e., the penalty for observation 1 is smaller (in magnitude) than the penalty for observation 2.

We present an example showing why the perpendicular distance from the end of the dashed line to the line $\mathbf{x}'\beta = 0$ does not completely determine the penalty, although it may seem at first that this might be the case. Consider two observations indexed by a and b with $\mathbf{x}'_a\beta = 5$ and $\exp(\mathbf{z}'_a\gamma) = 1$, and $\mathbf{x}'_b\beta = 25$ and $\exp(\mathbf{z}'_b\gamma) = 15$. For each of these two observations, as $\mathbf{x}'_i\beta > 0$ holds for $i = a, b$, the endpoint of the dashed line that is closest to the line $\mathbf{x}'\beta = 0$ is v_{i_1} in the notation of (7), and thus the distance from the end of each dashed lines to the line $\mathbf{x}'\beta = 0$ is determined by $\mathbf{x}'_i\beta - \exp(\mathbf{z}'_i\gamma)$. Set $u_i = \mathbf{x}'_i\beta - \exp(\mathbf{z}'_i\gamma)$. This distance is a strictly increasing function in u_i with value $d(u_i)$; i.e. it varies in an increasing fashion (albeit nonlinearly) with u_i . Now for observation a , this distance is $d(4) = d(5 - 1)$, whereas it is $d(10) = d(25 - 15)$ for observation b . The latter is larger, so one might be tempted given the way we have represented the model graphically to believe that the penalty would be smaller for observation b , as the dashed line is farther from the line $\mathbf{x}'\beta = 0$. But in fact, the ratio $\mathbf{x}'_i\beta / \exp(\mathbf{z}'_i\gamma)$ is larger for $i = b$ than it is for $i = 1$, as $15/25 > 1/5$, so the penalty for observation b is in fact larger. We could have drawn the dashed lines with length equal to the above ratio, but despite the flaws, we decided to draw the dashed lines as here, as it gives a representation of the range of values for the latent variable y^* , and because other alternatives seemed less natural and less apt to provide the reader with graphical intuition.

²⁵Here we have chosen $\gamma = (1, 5)$.




Figure 10: A Graphical Representation of a Plateau Solution in the Heteroskedastic Probit Model

plateau solution, the optimizer chooses γ with large and positive components.²⁶ The two planes for an illustrative case are presented in Figure 10. First look at the bottom panel. Observations with $\mathbf{z}_i = \mathbf{0}$, which we label $\mathbf{z}_1, \mathbf{z}_2$ and \mathbf{z}_3 for expositional convenience, of course lie directly on the line, while all of the other observations (represented here by unlabeled points) lie in the region $\mathbf{z}_i' \gamma > 0$ and away from the line. In the top graph, we see that only the points $\mathbf{x}_1, \mathbf{x}_2$ and \mathbf{x}_3 have small associated dashed line segments, while all others have large associated segments. (Note that some of these segments extend outside of the region of the graph.) So, by virtue of choosing γ with large, positive components, the optimizer has ensured that only the observations $i = 1, 2, 3$ matter very much for choosing the line $\mathbf{x}'\beta = 0$, i.e. for choosing β .

4 Potential Improvement and Other Suggestions

We now propose a straightforward method to mitigate the problems we have identified in the heteroskedastic probit model. We suggest ensuring that the model simply satisfies the property that at least one of the components of the vectors in \mathbf{z}_i take both signs. We would hope for this to avoid the plateau problem, or at least the exact form of it described above. Observe, however, that we do not claim that if one merely does so then maximizing the log likelihood function for the heteroskedastic probit model is robust; we are only stating that this is likely a helpful modification. Here we assume that there is no constant term in \mathbf{z} .²⁷ If there were a constant term in \mathbf{z}_i , say z_{i1} , one could simply choose the parameter γ_1 to be positive and very large relative to the other components of \mathbf{z}_i and thus ensure that $\exp(\mathbf{z}_i' \gamma)$ is always large, thus undoubtedly leading to a similar plateau problem.

We give a demonstration of the potential efficacy of our solution, by comparing the performance of an optimization algorithm on (i) a data set with $Z \geq 0$ and (ii) a translation of this first data set that ensures that the observations \mathbf{z}_i take both signs. The first data set

²⁶The choices we have made for these graphs, for ease of illustration, are $\beta = (-1.5, -5)$ and $\gamma = (8, 2)$.

²⁷There may or may not be a constant term in \mathbf{x} here. In general, there should of course not be a constant term in *both* \mathbf{x} and \mathbf{z} , as the model would not be identified.

is created as described above, and, for example, is forced to satisfy the requirement that the percentage of observations with crossover is between 20% and 30%. Denote the data set by \mathcal{D} . Let n be the number of observations in the data set. \mathcal{D} consists of the parameters β_0 and γ_0 and the observations \mathbf{y} , X and Z , so we write

$$\mathcal{D} = \{\beta_0, \gamma_0, \mathbf{y}, X, Z\}.$$

Also define $k_1 = \dim(\beta_0)$ and $k_2 = \dim(\gamma_0)$. Observe that we have $\dim(\mathbf{y}) = n$, while $\dim(X) = n \times k_1$ and $\dim(Z) = n \times k_2$. Write

$$\gamma_0 = (\gamma_{01}, \gamma_{02}, \dots, \gamma_{0k_2}).$$

For convenience, write

$$\mathbf{1} = \underbrace{(1, 1, \dots, 1)}_{\text{length } k_2}.$$

Now we define the transformed data set $\tilde{\mathcal{D}}$. Let

$$\begin{aligned} \tilde{\beta}_0 &= \left(\exp \left(- \sum_{j=1}^{k_2} \frac{\gamma_{0j}}{2} \right) \right) \beta_0, \\ \tilde{\gamma}_0 &= \gamma_0, \\ \tilde{\mathbf{y}} &= \mathbf{y}, \\ \tilde{\mathbf{x}}_i &= \mathbf{x}_i \quad \text{for } 1 \leq i \leq n, \\ \tilde{\mathbf{z}}_i &= \mathbf{z}_i - \frac{1}{2} \mathbf{1} = \mathbf{z}_i - \frac{1}{2} \underbrace{(1, 1, \dots, 1)}_{\text{length } k_2} \quad \text{for } 1 \leq i \leq n, \end{aligned}$$

and

$$\tilde{\mathcal{D}} = \{\tilde{\beta}_0, \tilde{\gamma}_0, \tilde{\mathbf{y}}, \tilde{\mathbf{x}}, \tilde{\mathbf{z}}\}.$$

We note for emphasis that we have $\dim(\tilde{\beta}_0) = k_1$ and $\dim(\tilde{\gamma}_0) = k_2$, and also $\dim(\tilde{\mathbf{y}}) = n$, while $\dim(\tilde{X}) = n \times k_1$ and $\dim(\tilde{Z}) = n \times k_2$.

Now, consider, for each i , the expressions $\frac{\mathbf{x}'_i \boldsymbol{\beta}}{\exp(\mathbf{z}'_i \boldsymbol{\gamma})}$ and $\frac{\tilde{\mathbf{x}}'_i \tilde{\boldsymbol{\beta}}}{\exp(\tilde{\mathbf{z}}'_i \tilde{\boldsymbol{\gamma}})}$. These appear as terms in each summand of the respective log likelihood functions for the heteroskedastic probit model, for the original and transformed data sets. We claim that for all i , the value of this term for the original data set equals its value for the transformed data set at each data set's respective model parameters: for we have

$$\begin{aligned} \frac{\tilde{\mathbf{x}}'_i \tilde{\boldsymbol{\beta}}_0}{\exp(\tilde{\mathbf{z}}'_i \tilde{\boldsymbol{\gamma}}_0)} &= \left(\exp \left(- \sum_{j=1}^{k_2} \frac{\gamma_{0j}}{2} \right) \right) \frac{\mathbf{x}'_i \boldsymbol{\beta}_0}{\exp \left((\mathbf{z}_i - \frac{1}{2} \mathbf{1})' \boldsymbol{\gamma}_0 \right)} \\ &= \left(\exp \left(- \sum_{j=1}^{k_2} \frac{\gamma_{0j}}{2} \right) \right) \frac{\mathbf{x}'_i \boldsymbol{\beta}_0}{\exp(\mathbf{z}'_i \boldsymbol{\gamma}_0) \exp(-\frac{1}{2} \mathbf{1}' \boldsymbol{\gamma}_0)} \\ &= \frac{\mathbf{x}'_i \boldsymbol{\beta}_0}{\exp(\mathbf{z}'_i \boldsymbol{\gamma}_0)}. \end{aligned}$$

Together with the observation that $\tilde{\mathbf{y}} = \mathbf{y}$, we can see that the i^{th} summand of the log likelihood function for the original data set at its model parameters $(\boldsymbol{\beta}_0, \boldsymbol{\gamma}_0)$ is equal to the i^{th} summand of the log likelihood function for the transformed data set at its model parameters $(\tilde{\boldsymbol{\beta}}_0, \tilde{\boldsymbol{\gamma}}_0)$. In other words, writing $\mathbf{y} = (y_1, y_2, \dots, y_n)$ and $\tilde{\mathbf{y}} = (\tilde{y}_1, \tilde{y}_2, \dots, \tilde{y}_n)$, we have:

$$\begin{aligned} \left(y_i \ln \Phi \left(\frac{\mathbf{x}'_i \boldsymbol{\beta}}{\exp(\mathbf{z}'_i \boldsymbol{\gamma})} \right) + (1 - y_i) \ln \left(1 - \Phi \left(\frac{\mathbf{x}'_i \boldsymbol{\beta}}{\exp(\mathbf{z}'_i \boldsymbol{\gamma})} \right) \right) \right) &= \\ \left(\tilde{y}_i \ln \Phi \left(\frac{\tilde{\mathbf{x}}'_i \tilde{\boldsymbol{\beta}}_0}{\exp(\tilde{\mathbf{z}}'_i \tilde{\boldsymbol{\gamma}}_0)} \right) + (1 - \tilde{y}_i) \ln \left(1 - \Phi \left(\frac{\tilde{\mathbf{x}}'_i \tilde{\boldsymbol{\beta}}_0}{\exp(\tilde{\mathbf{z}}'_i \tilde{\boldsymbol{\gamma}}_0)} \right) \right) \right) & \\ \text{for } 1 \leq i \leq n. & \end{aligned}$$

It of course follows immediately that the respective log likelihood functions at the respective model parameters are equal, i.e. that

$$\ell(\boldsymbol{\beta}_0, \boldsymbol{\gamma}_0 \mid y, X, Z) = \ell(\tilde{\boldsymbol{\beta}}_0, \tilde{\boldsymbol{\gamma}}_0 \mid \tilde{y}, \tilde{X}, \tilde{Z}).$$

Now we turn to some details. As described above, we create a pseudorandom set \mathcal{D} . We then generate 1000 random choices of initial values in the box $[-5, 5]^5$, and measure the search performance of the BFGS algorithm in estimating the model parameters (β_0, γ_0) . Then we transform the data set \mathcal{D} as described above to get a new data set $\tilde{\mathcal{D}}$. We then transform these 1000 new random choices of initial values (in the same way we transformed β in the equations above, keeping γ unchanged), and subsequently measure the search performance of the BFGS algorithm in approaching its model parameters $(\tilde{\beta}_0, \tilde{\gamma}_0)$.²⁸

We present the results in Figure 11. First consider the left graphs. The top left plot displays a histogram of the Euclidean distances between the estimated and model parameter values for the 1000 runs for the original data set \mathcal{D} ; the bottom left plot displays the analogous histogram for the transformed data set $\tilde{\mathcal{D}}$.²⁹ Clearly the performance in the latter case is far superior. Now consider the right column of graphs. As with earlier plots, we also display histograms of the (signed) differences between the normalized values of the log likelihood at the model and estimated parameters. (Recall that the former might generally be expected to be greater, i.e. a smaller negative number.) The top graph corresponds to the original simulated data set; the bottom plot to the transformed data set. Horizontal and vertical axes in both graphs have the same scale. Again, in both graphs, we have also plotted a bar that represents the number of choices of initial values for which the log likelihood value at the resulting estimates is actually larger than the log likelihood at the model parameters.³⁰ There are many more of these for the transformed data set, and in fact for the transformed data set, there are only three choices of initial values (out of 1000) for which the log likelihood at the estimated parameters is worse than (i.e. less than) the log likelihood at the model parameters. Clearly, by this method of comparison as well, the BFGS search algorithm performs much better on the transformed data set.

Naturally, we have only presented evidence for one simulated data set and one search

²⁸We also ran a simulation on the same data set but where we chose different initial values for each of the original and transformed data set. The results are very similar.

²⁹Observe that the plots have the same scales on the horizontal and vertical axes.

³⁰As above, the horizontal position of this bar is of course not meaningful.

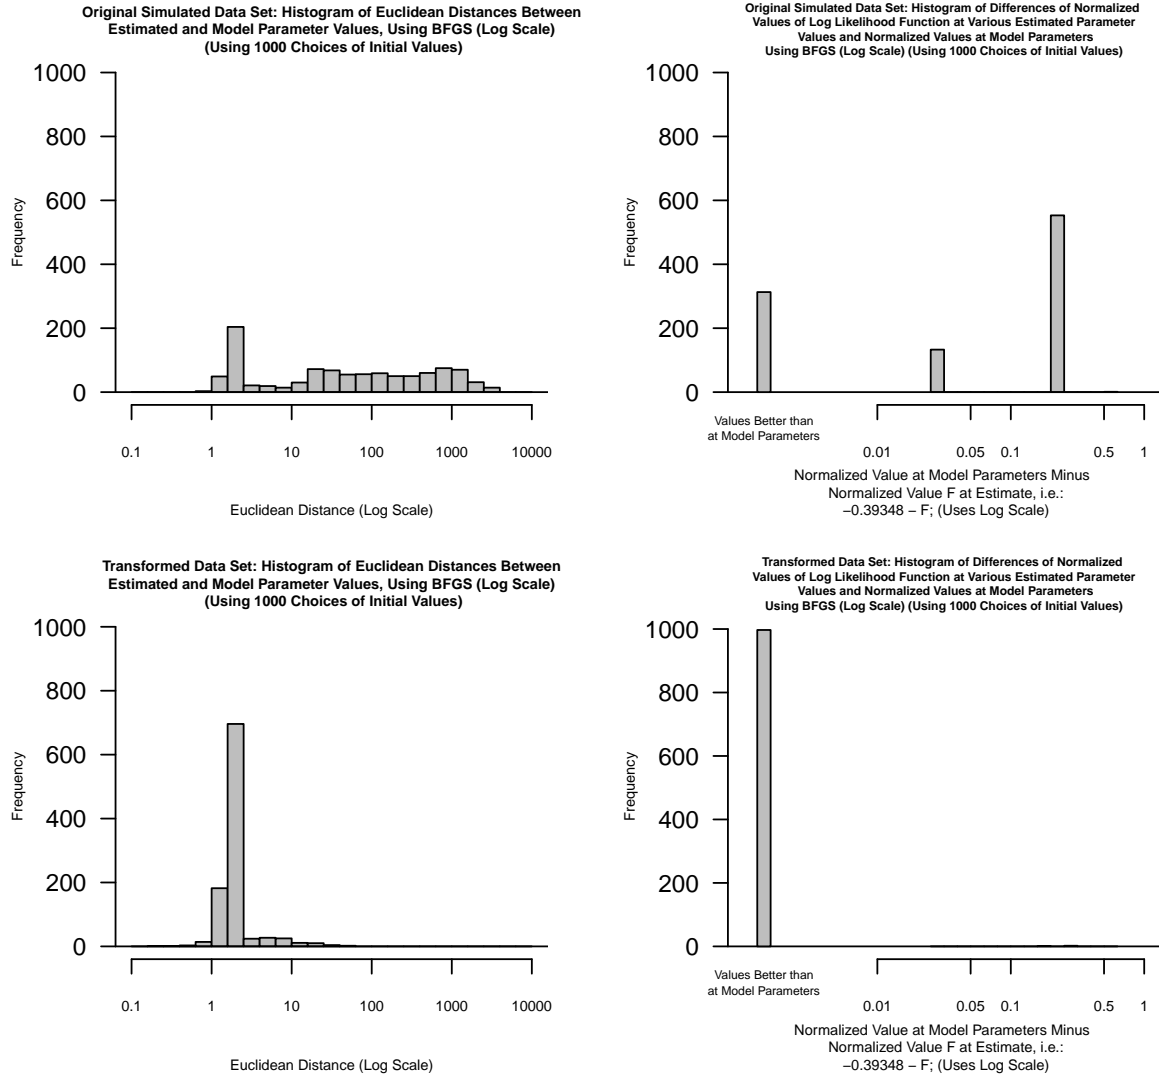


Figure 11: Comparison of Performance of the BFGS Search Algorithm for a Simulated Data Set and a Transformation of that Data Set (Using the Heteroskedastic Probit Model), for 1000 Random Choices of Initial Values

algorithm.³¹ But we have given reasons why this transformation may mitigate the plateau problem.³²

³¹We also performed similar simulations for the CG, Nelder-Mead and SANN algorithms. The CG algorithm performs much better for the transformed data set than for the original data set, and the Nelder-Mead method performs better but not strongly so; the SANN algorithm on the other hand performs fairly well for both data sets.

³²It may be of interest to discuss, in terms of a search algorithm, how the transformation affects the plateau problem. Note that parameter values on the plateau for the original set, i.e. with large values of γ , correspond, under the transformation, to values of β close to $\mathbf{0}$; this can be seen from the definition of the

4.1 Additional Suggestions

In addition to the above modification, we make a few other suggestions that may be helpful for researchers to bear in mind when using a search algorithm to estimate a heteroskedastic probit model.

One thing we suggest is that a researcher should try multiple different starting values for search algorithms. One could then check if one gets a stable set of resulting parameter estimates; if not, caution is called for. This is presumably a good approach in general, but it appears to be especially warranted for the heteroskedastic probit model. As well, if a researcher uses a search algorithm, and finds an estimate $\hat{\gamma}$ that has large positive components, there should be concern that it is a plateau solution; in line with the previous suggestion, the researcher should try many other initial values. If that does not improve the situation, results should be interpreted with caution. The simulated annealing algorithm (SANN in this implementation) is, in some respects, similar in spirit to the notion of starting search algorithms at multiple initial values, in that it incorporates random jumps in the procedure. In some of the simulations SANN performed much better than some of the other algorithms; however, it did not perform as well as other methods on the Alvarez-Brehm data. So a simulated annealing algorithm might be a useful search method for the heteroskedastic probit. We emphasize that we do not view a simulated annealing algorithm as a panacea by any means, and we are not implying that simply using a simulated annealing method suffices to address the problems we have identified in the heteroskedastic probit model.

Finally, we make an observation regarding the special case in which \mathbf{x} and \mathbf{z} are independent. From (4), we can see that intuitively, for positive and large choices of $\hat{\gamma}$, it seems that the choice of $\hat{\beta}$ would be determined by the optimal choice of $\hat{\beta}$ for the subset of i for which $\mathbf{z}_i = \mathbf{0}$. If \mathbf{x} and \mathbf{z} are independent, then the observations \mathbf{x}_i for i in this subset are a random sample, and thus we would expect the estimate of β_0 to be good in some sense if the estimate $\hat{\gamma}$ is positive and large. Thus this presents a potential (informal) method of

transformation. Thus, small jumps in a local search on the transformed set would lead a search algorithm away from $\mathbf{0}$, which would correspond to escaping the plateau in the parameter space for the original data set.

estimation for the heteroskedastic probit in this special case: first, find an estimate $\hat{\beta}$ and if the associated estimate $\hat{\gamma}$ is positive and large, then, second, fix this value of $\hat{\beta}$ and use a search algorithm to estimate only the γ_0 parameter.³³

5 Final Remarks

Here we have explored how commonly used optimization algorithms perform when applied to the heteroskedastic probit likelihood. Using both simulated data sets and in a re-analysis of the seminal work by Alvarez and Brehm, we find that some optimization algorithms can often converge at incorrect solutions. We compare these same algorithms when applied to standard probit log likelihood functions and find no such problems. We also sketch a heuristic argument for why these difficulties may be inherent in the heteroskedastic probit likelihood. Our work constitutes a first step, by calling attention to some potential problems in using the model. But more remains to be done to understand the particular functional form. However, analysts that estimate heteroskedastic probit models should at the very least ensure that a variety of starting values converge to the same set of estimates. They may also want to perform a transformation of the data and parameters in the fashion we have described. It would also be advisable to try more robust optimization algorithms such as simulated annealing or a genetic optimization algorithm (Sekhon and Mebane 1998).

³³An alternative method might involve starting with a positive and large initial value for γ . Also, we have observed the estimate of β_0 to be close to the model or A-B parameters in multiple simulations. But we do not investigate making either of the proposed methods more formal, as this is not the focus of this paper.

References

- Albert, A and J. A. Anderson. 1984. "On the Existence of Maximum Likelihood Estimates in Logistic Regression Models." *Biometrika* 71(1):1–10.
- Alvarez, R. Michael and John Brehm. 1995. "American Ambivalence Towards Abortion Policy: Development of a Heteroskedastic Probit Model of Competing Values." *American Journal of Political Science* 39(4):1055–1082.
- Alvarez, R. Michael and John Brehm. 1997. "Are Americans Ambivalent Towards Racial Policies." *American Journal of Political Science* 41(2):345–374.
- Alvarez, R. Michael and John Brehm. 1998. "Speaking In Two Voices: American Equivocation About the Internal Revenue Service." *American Journal of Political Science* 42(2):418–452.
- Alvarez, R. Michael and John Brehm. 2002. *Hard Choices, Easy Answers*. Princeton, NJ: Princeton University Press.
- Belisle, C. J. P. 1992. "Convergence Theorems for a Class of Simulated Annealing Algorithms on \mathbb{R}^d ." *Journal of Applied Probability* 29:885–895.
- Broyden, C.G. 1970. "The Convergence of a Class of Double-Rank Minimization Algorithms." *Journal of the Institute of Mathematics and Its Application* 6:76–90.
- Busch, Marc L. and Eric Reinhardt. 1999. "Industrial Location and Protection: The Political and Economic Geography of U.S. Nontariff Barriers." *American Journal of Political Science* 43(4):1028–1050.
- Fletcher, R. 1970. "A New Approach to Variable Metric Algorithms." *Computer Journal* 13:317–322.
- Fletcher, R and C. M. Reeves. 1964. "Function Minimization by Conjugate Gradients." *Computer Journal* 7:148–154.

- Gabel, M. J. 1998. "Economic Integration and Mass Politics: Market Liberalization and Public Attitudes in the European Union." *American Journal of Political Science* 42(3):936–953.
- Harvey, Andrew. 1976. "Estimating Regression Models with Multiplicative Heteroskedasticity." *Econometrica* 44:461–465.
- Iyengar, Satish and Joel B. Greenhouse. 1988. "Selection Models and the File Drawer Problem." *Statistical Science* 3(1):109–135.
- Krutz, Glen S. 2005. "Issues and Institutions: Winnowing in the U.S. Congress." *American Journal of Political Science* 49:436–447.
- Lee, Taeku. 2002. *Mobilizing Public Opinion: Black Insurgency and Racial Attitudes in the Civil Rights Era*. Chicago: University of Chicago Press.
- Nelder, J.A. and R. Mead. 1965. "A Simplex Algorithm for Function Minimization." *Comparative Journal* 7:308–313.
- Ruud, Paul, A. 2000. *An Introduction to Classical Econometric Theory*. Oxford: Oxford University Press.
- Sekhon, Jasjeet S. and Walter R. Mebane. 1998. "Genetic Optimization Using Derivative: Theory and Application to Nonlinear Models." *Political Analysis* 7:189–213.
- Shanno, David F. 1970. "Conditioning of quasi-Newton Methods for Function Minimization." *Mathematics of Computation* 24:647–656.
- Yatchew, A. and Z. Griliches. 1985. "Specification Error In Probit Models." *Review of Economics and Statistics* 18:134–139.

A Appendix: Graphs for the SANN Search Algorithm

Figure 12: Performance of the SANN Search Algorithm on a Simulated Data Set for the Heteroskedastic Probit Model, for 1000 Random Choices of Initial Values




Figure 13: Performance of the SANN Search Algorithm on a Simulated Data Set for the Probit Model, for 1000 Random Choices of Initial Values

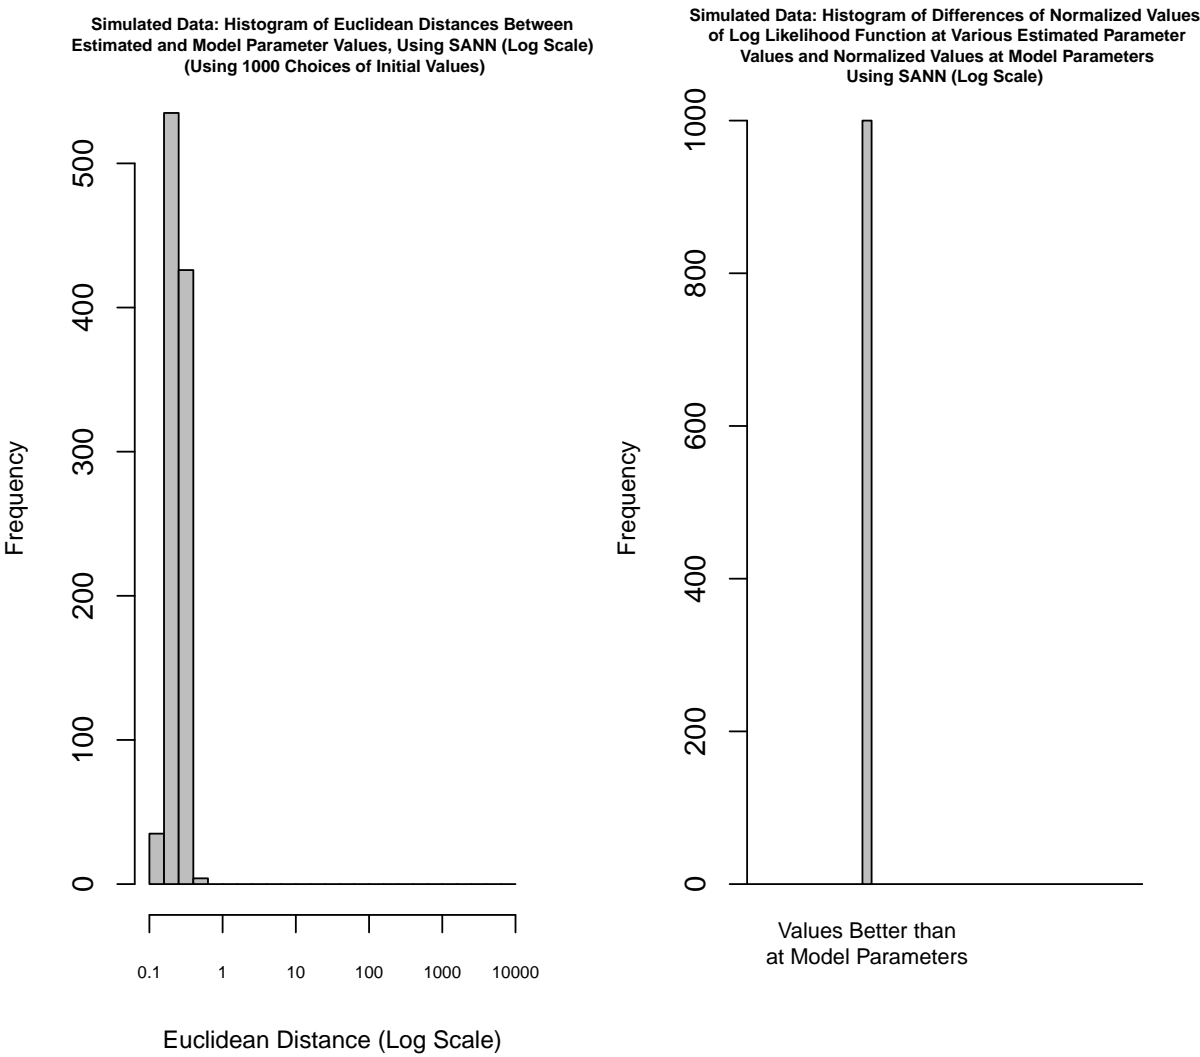


Figure 14: Performance of the SANN Search Algorithm on the Alvarez-Brehm Data, for 1000 Random Choices of Initial Values

